

RECEIVED  
CENTRAL FAX CENTER  
NOV 03 2006

Amendments to the Specification

Replace paragraph [0046] with the following:

FIG. 4A depicts a storage chapter table. As noted above, a chapter is a plurality of paragraphs (rows or sheets) to which all sparing versions have been consecutively applied. Chapter table 400 contains paragraphs 402 and drives designated as group 1 (404), group 2 (406), or spare 408. Letters in each row may represent a single data block on that drive, or may represent a plurality of contiguous storage blocks, called a super block, on that drive. For example, if the paragraph demarcation is a sheet, and the sheet comprises a predefined number of rows, then the number of data blocks in a super block is equal to the number of rows in the sheet. It may be observed from chapter table 400 that a uniform amount of user storage capacity (i.e. excluding spares) exists in each paragraph. The group of user data blocks (that may be used for data or parity) in the plurality of rows of a sheet to which sparing has been applied is termed a capacity grid. In other words, a sheet organization is defined for a plurality of sheets, sparing (if any) is applied to the sheets, resulting in a plurality of capacity grids. FIG. 4B depicts a capacity grid. Capacity grid 410 is a grid that may be produced by defining sheets having five rows that are spared at a sheet boundary (i.e. the paragraph comprises five rows) where the first sparing version is employed. The drives providing user data in each group are the same within each row of the paragraph, as illustrated by A capacity grid may be viewed as the available storage of a sheet to which sparing, pairing, and grouping are transparently applied. The capacity grid 410 contains 2 groups, group 1 (412) comprising drives A, C, E, and G and group 2 (414) comprising drives B, D, F, and H. The data blocks in a row of the capacity grid that belong to one group (four data blocks in this example) are termed a data storage unit (DSU) and as such ten DSUs are

depicted in capacity grid 410. DSUs may be considered as belonging to group 1 or group 2. In some embodiments of the present invention, all DSUs in a capacity grid are allocated to the same logical device, reducing virtualization metadata overhead.

Replace paragraphs [0053-0054] with the following:

FIG. 10A depicts a RAID-6 capacity grid. RAID-6 is similar to RAID-5 in that a row parity value is calculated for each DSU but also includes a vertical parity that employs XORing data blocks from different DSUs. Grid 1000 comprises eight rows containing sixteen data ~~stets~~ sets and four rows containing parity. Data sets (four data blocks each) are represented by the letters A through P.

FIG. 10B depicts a RAID-6 DSU/Parity Map. One embodiment of RAID-6 parity calculations may be understood through RAID-6 DSU/Parity Map 1002 shown in FIG. 10B. Whereas grid 1000 depicts ~~datasets~~ data sets and parity rows, DSU/Parity map 1002 shows data blocks 1004 that contain user data and parity blocks 1006 that contain row or vertical parity values. Each data block contains a letter-number index pair. The letter index designates which data blocks are employed to calculate a row parity value and the number index designates which blocks are employed to calculate a vertical parity value. Data blocks having the same index are XOR'ed (logically exclusive OR'ed) to produce a parity value. For example, A-1, A-2, A-3 and A-4 are XORed to produce AP and A-1, B-1, C-1 and D-1 are XORed to produce 1P. Each data block contributes to one row parity value and to one vertical parity value.